Eye mask measurement accuracy

Piers Dawe

For IEEE P802.3ah EFM, May 2003, Seoul

Agilent Technologies

Problem statements:

1. Eye mask margin measurements can be misleading and less than accurate

2. Eye measurement procedure is not fully specified in reference standards

3. Seeking correlation between TDP and other measurement methods

This is a general analysis, not specific to any bit rate

Eye mask addresses high probability impairments ("deterministic"); various noise effects can give misleading results

Mask measurement: Scope and Tx noises means measuring for longer makes apparent eye smaller

Tails of distributions do not affect link performance: they are drowned by link's Rx noise and jitter (noises combine as sum of squares)

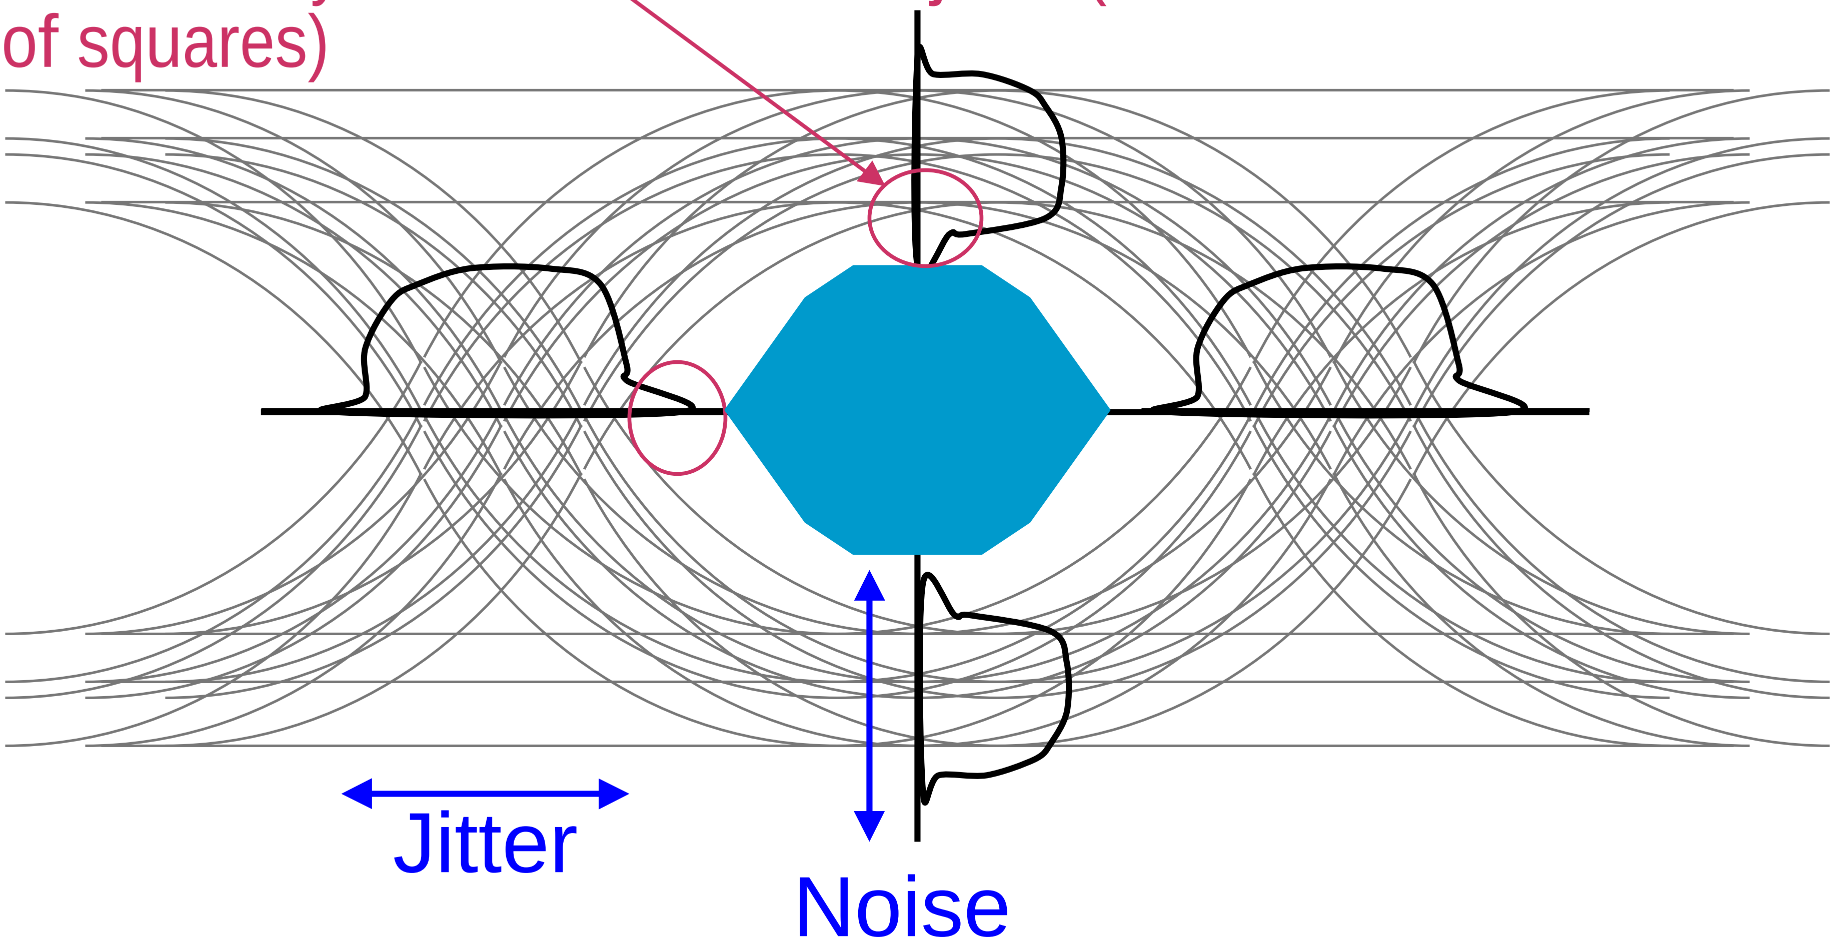

Noises combine as sum of squares: Typical example for 10G Ethernet measurement

Probability distribution (Tx under test) combines with probability distribution (scope) to determine mask hit rate

(Measurement technique determines scatter in measurement)

Measurement:

$(\text{Tx noise})^2 + (\text{scope noise})^2 = (\text{measurement total noise})^2$

example: $0.03^2 + 0.02^2 = 0.036^2$

Target use:

$(\text{Tx noise})^2 + (\text{Rx noise})^2 = (\text{system total noise})^2$

example: $0.03^2 + 0.139^2 = 0.142^2$

Drowned by link Rx, not important at this level

And see backup

Scope and Tx noises means measuring for longer makes apparent eye smaller:

Distributions of both DUT and instrument contribute to mask hits

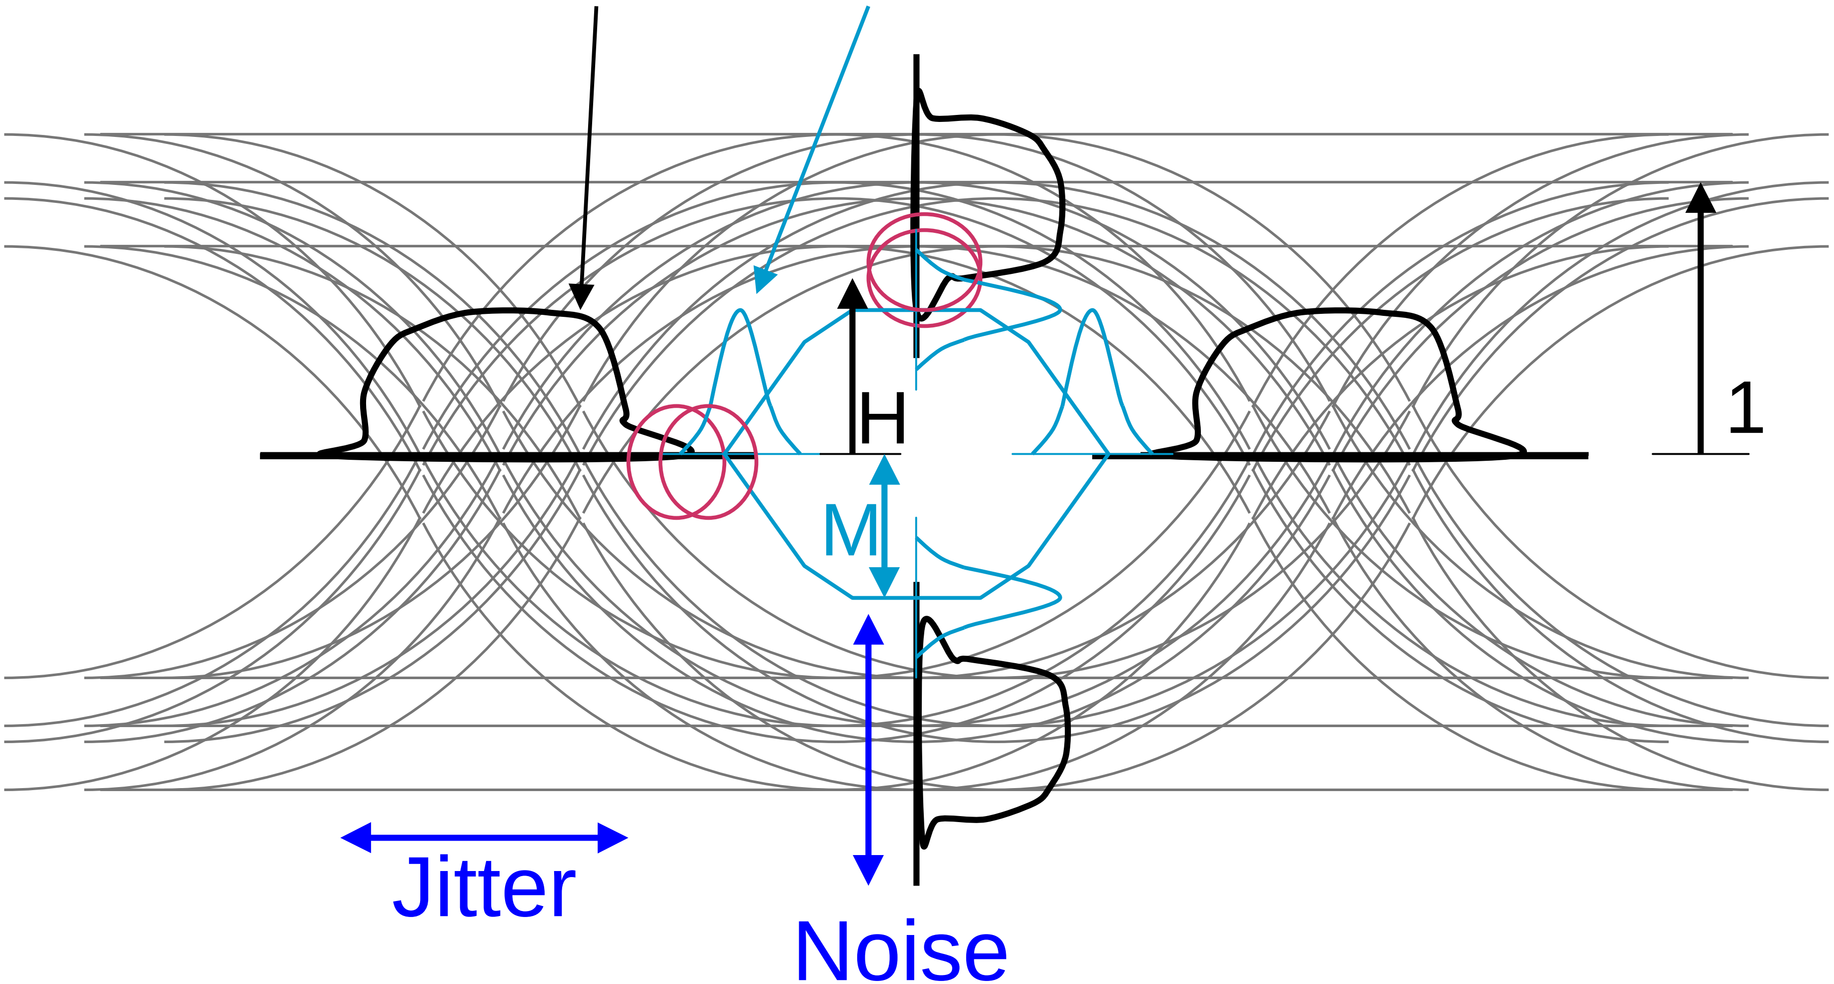

Real receiver noise is much greater than scope or DUT noise

In actual use, tails of Tx distributions are drowned by link’s Rx noise and jitter

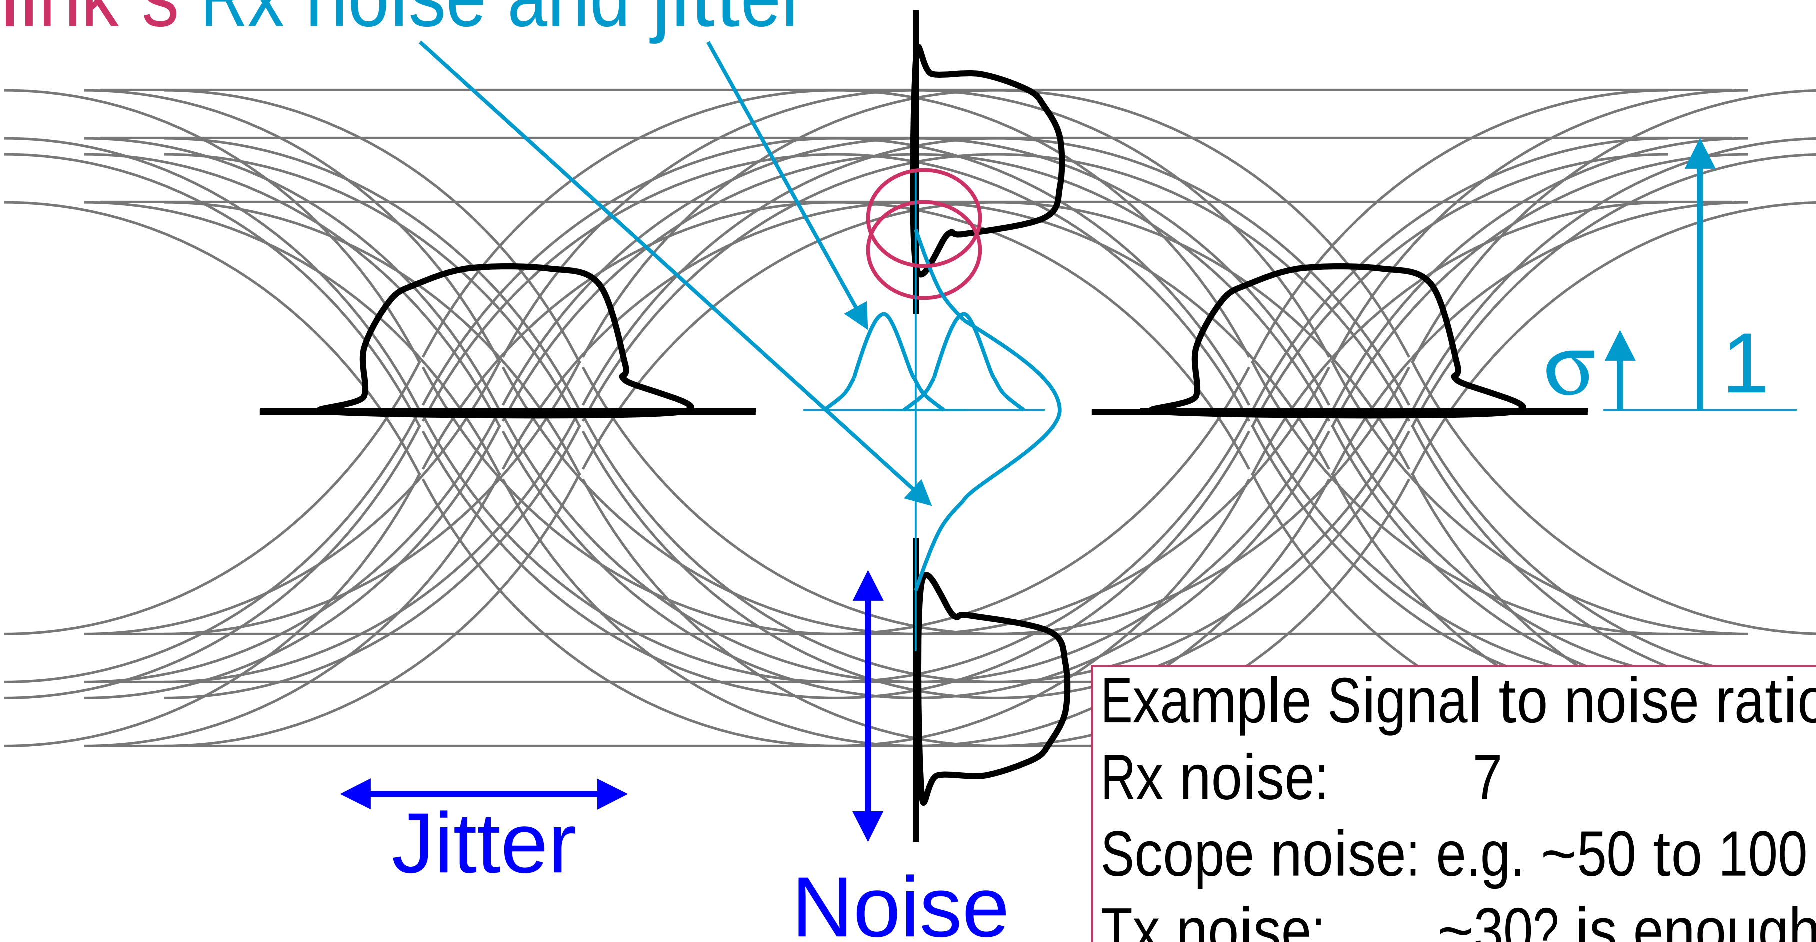

Example Signal to noise ratios:

Rx noise: 7

Scope noise: e.g. ~50 to 100

Tx noise: ~30? is enough

Number of samples

Even modern sampling scopes miss most of the signal

e.g. 40 ksamples / s

as slide 5

Good for finding high probability effects and for diagnostic analysis

Actual system or TDP test takes 125,000 to 10,000,000 ksamples / s

as slide 6

Do not recommend using eye mask on scope (DCA) to characterise low probabilities:
TDP spec addresses these

Theoretical analysis of relation between actual transmitter penalty and observed mask margin: method

1. Calculate relation between transmitter penalty (TP) and mask margin (MM) assuming no noise anywhere:

TP = 10*log10(1/H), MM=(H-M)/(1-M) where H is height of inner eye and M is height of mask*

2. Extend calculation to allow for penalty of Tx noise
3. Extend calculation to include scope noise
4. Find position in tail of distribution representing 0 hits
5. Find likely scatter around that point
6. Graph out likely range of apparent mask margin vs. actual transmitter penalty for a family of transmitters from good to bad

* See slide 5   Analysis assumes =0.5

Theoretical analysis of relation between actual transmitter penalty and observed mask margin

Basing mask margin on zero hits may produce widely varying results!

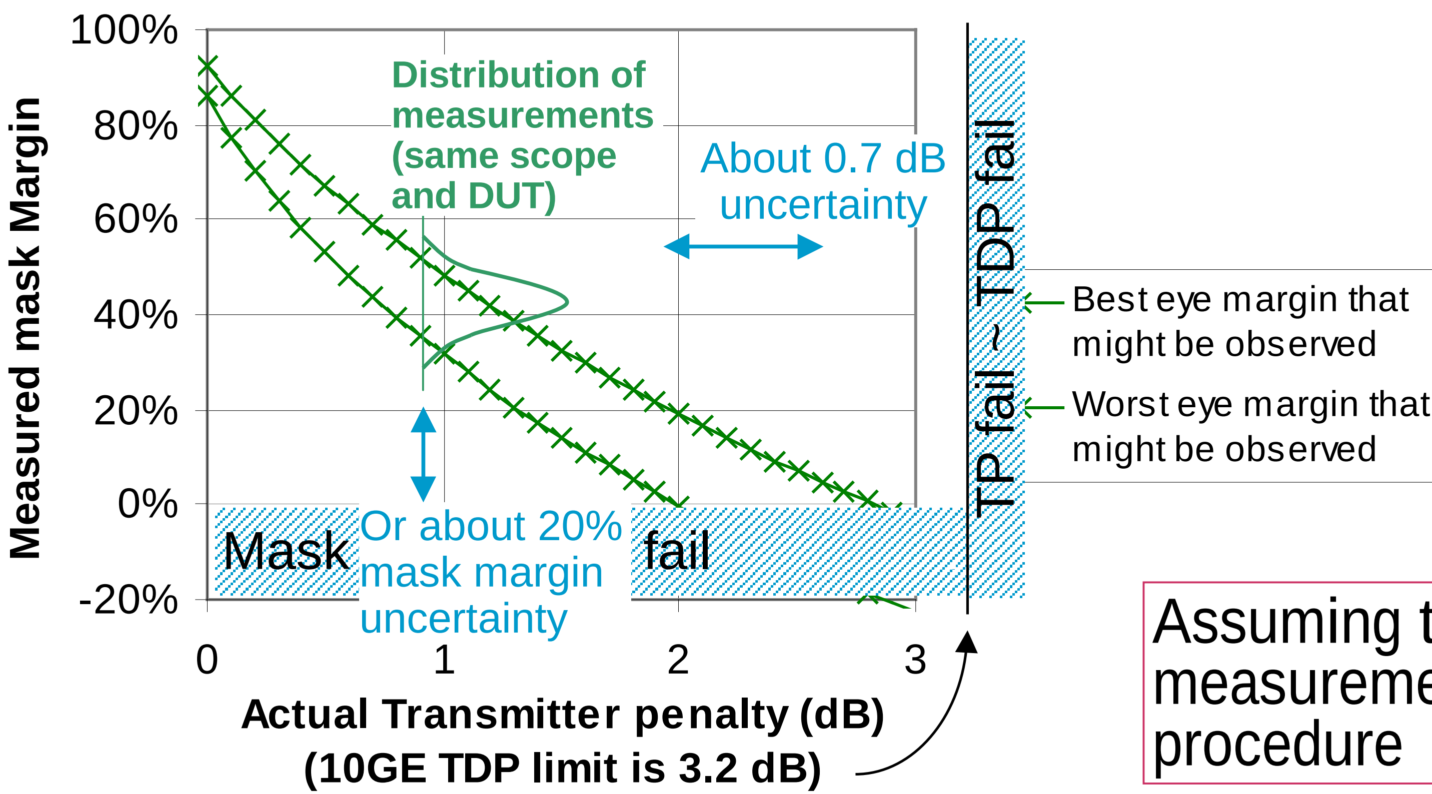

Assuming typical measurement procedure

Theoretical analysis of relation between actual transmitter penalty and observed mask margin

Basing mask margin on zero hits may produce widely varying results!

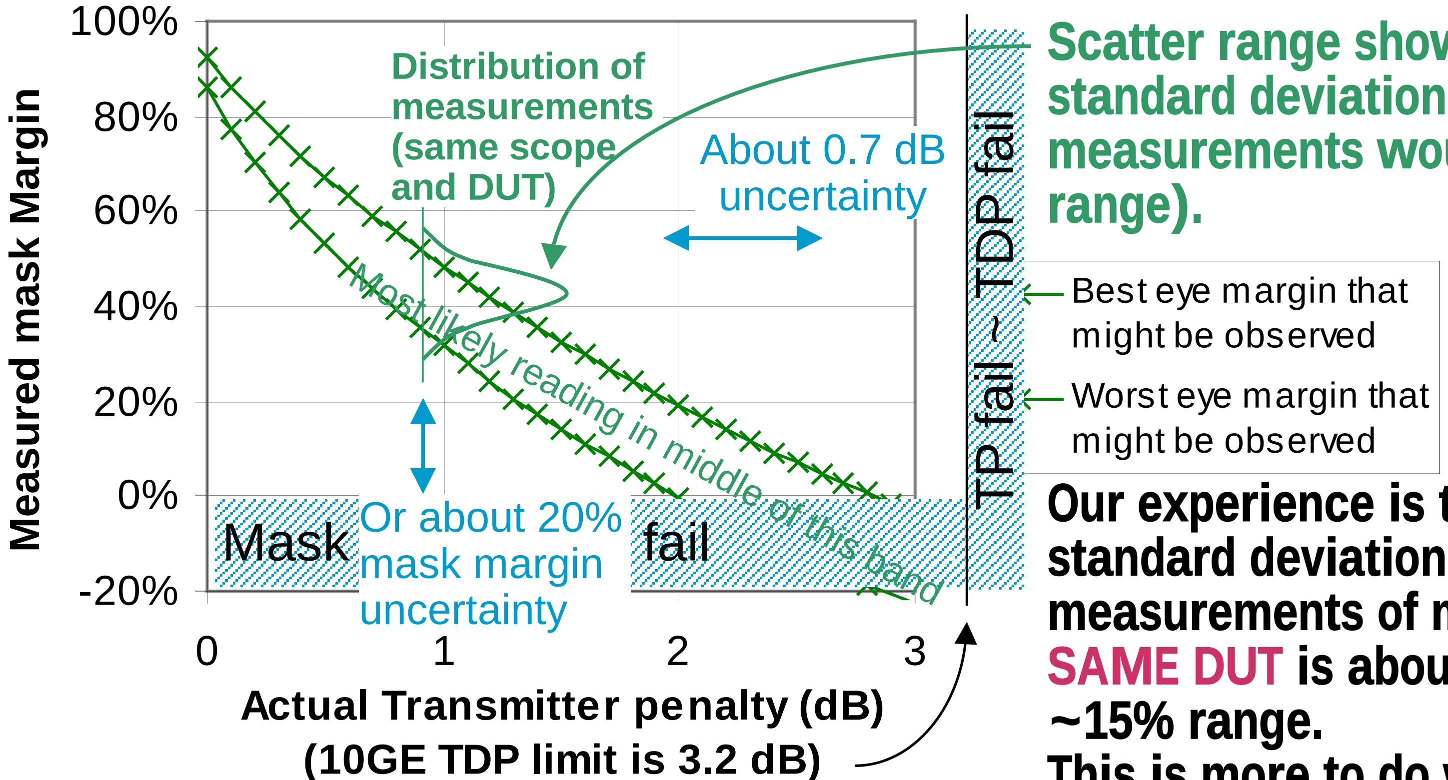

Scatter range shown is +-2.5 standard deviations (99% of measurements would fall in this range).

Our experience is that the standard deviation of repeated measurements of mask margin ON SAME DUT is about 3%, giving ~15% range.
This is more to do with the measurement than the DUT.
New lower noise scope plug-in helps.

Industry standard: 200 waveforms, no hits

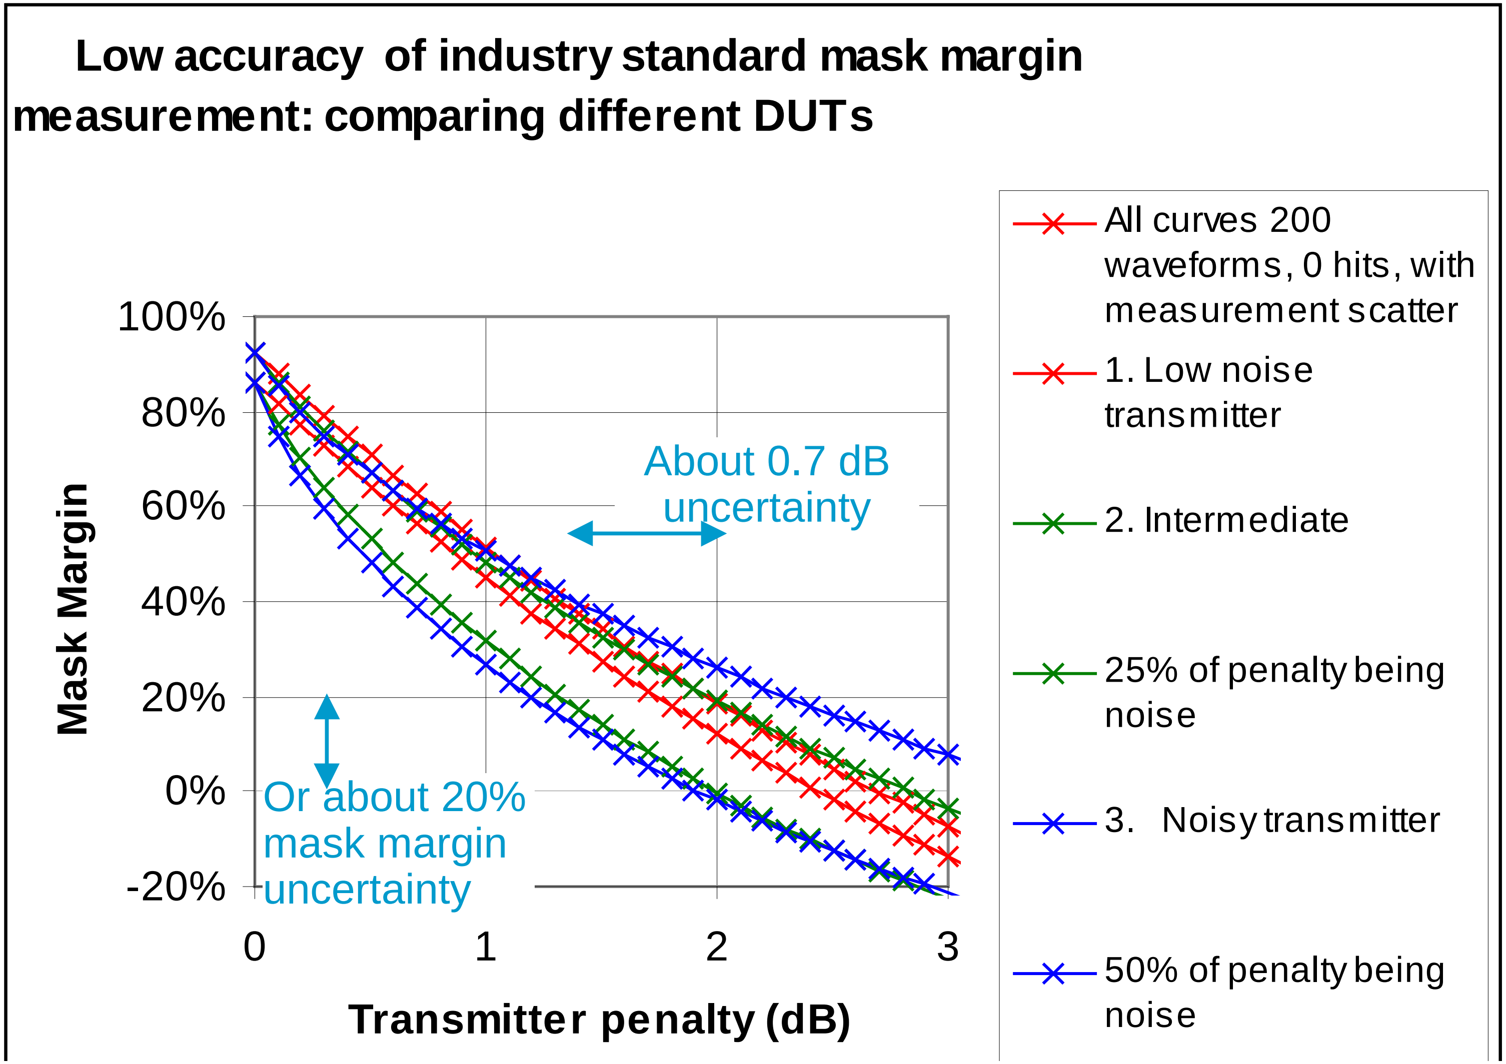

Now showing transmitters with different types of impairment: deterministic and random. 3 examples in 3 colours.

Can we get better accuracy with more waveforms?

No

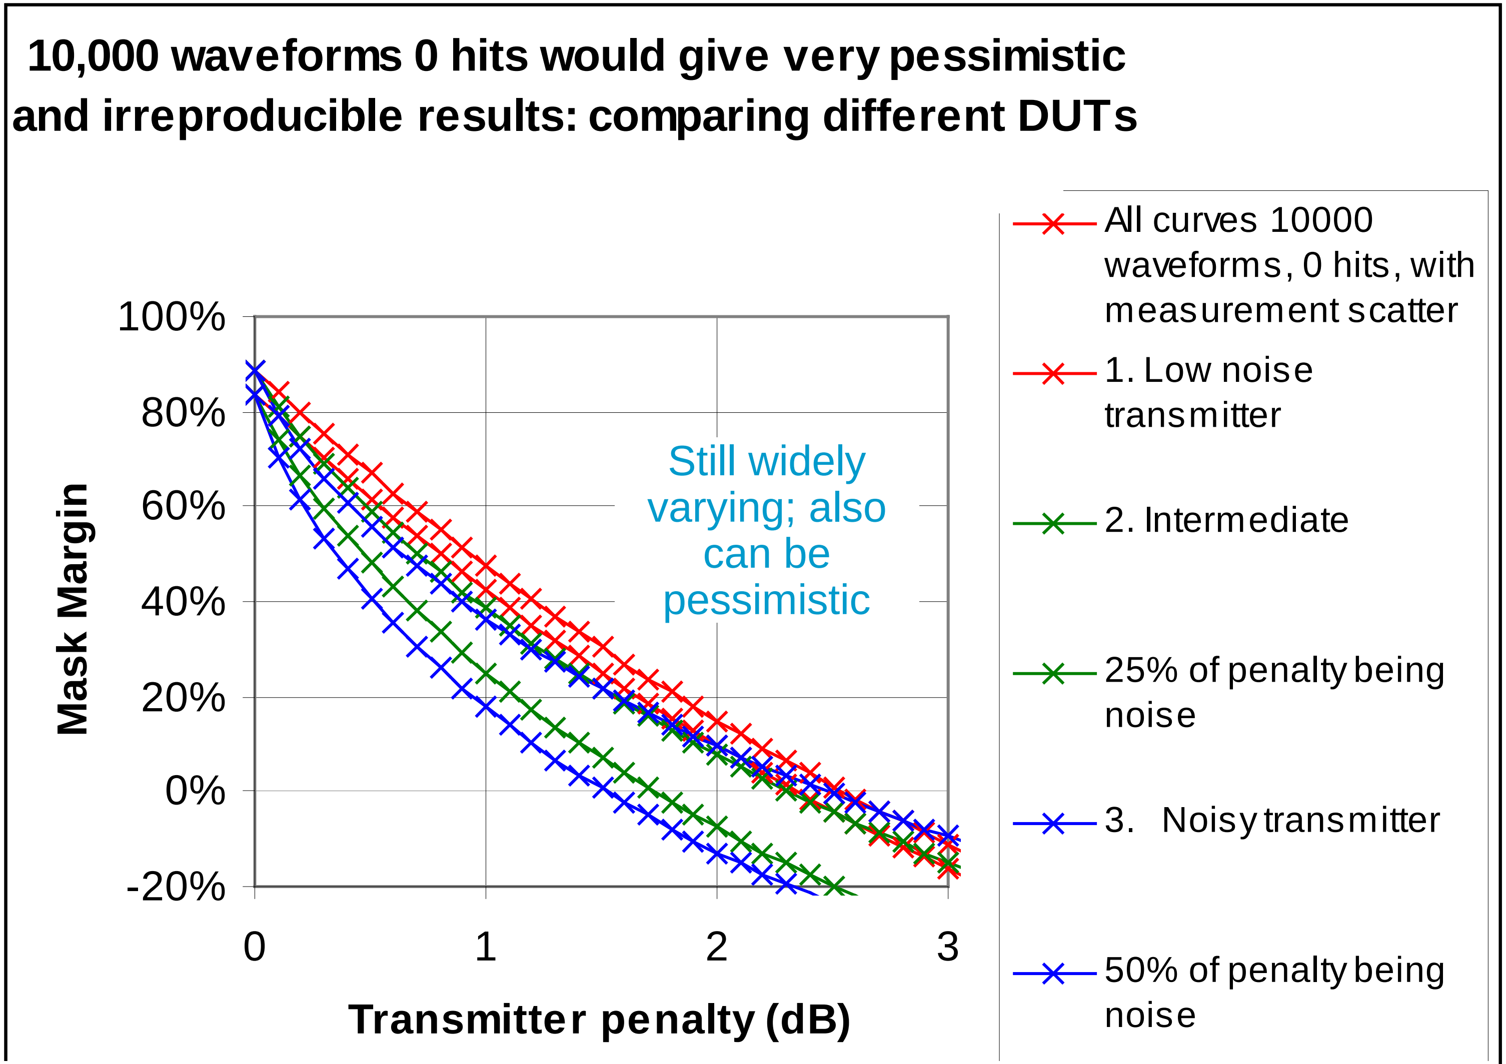

Note some of these curves are more pessimistic than previous page ("200 waveforms") by up to 10%

Intermediate position still suffers from over-measured noise and scattered results

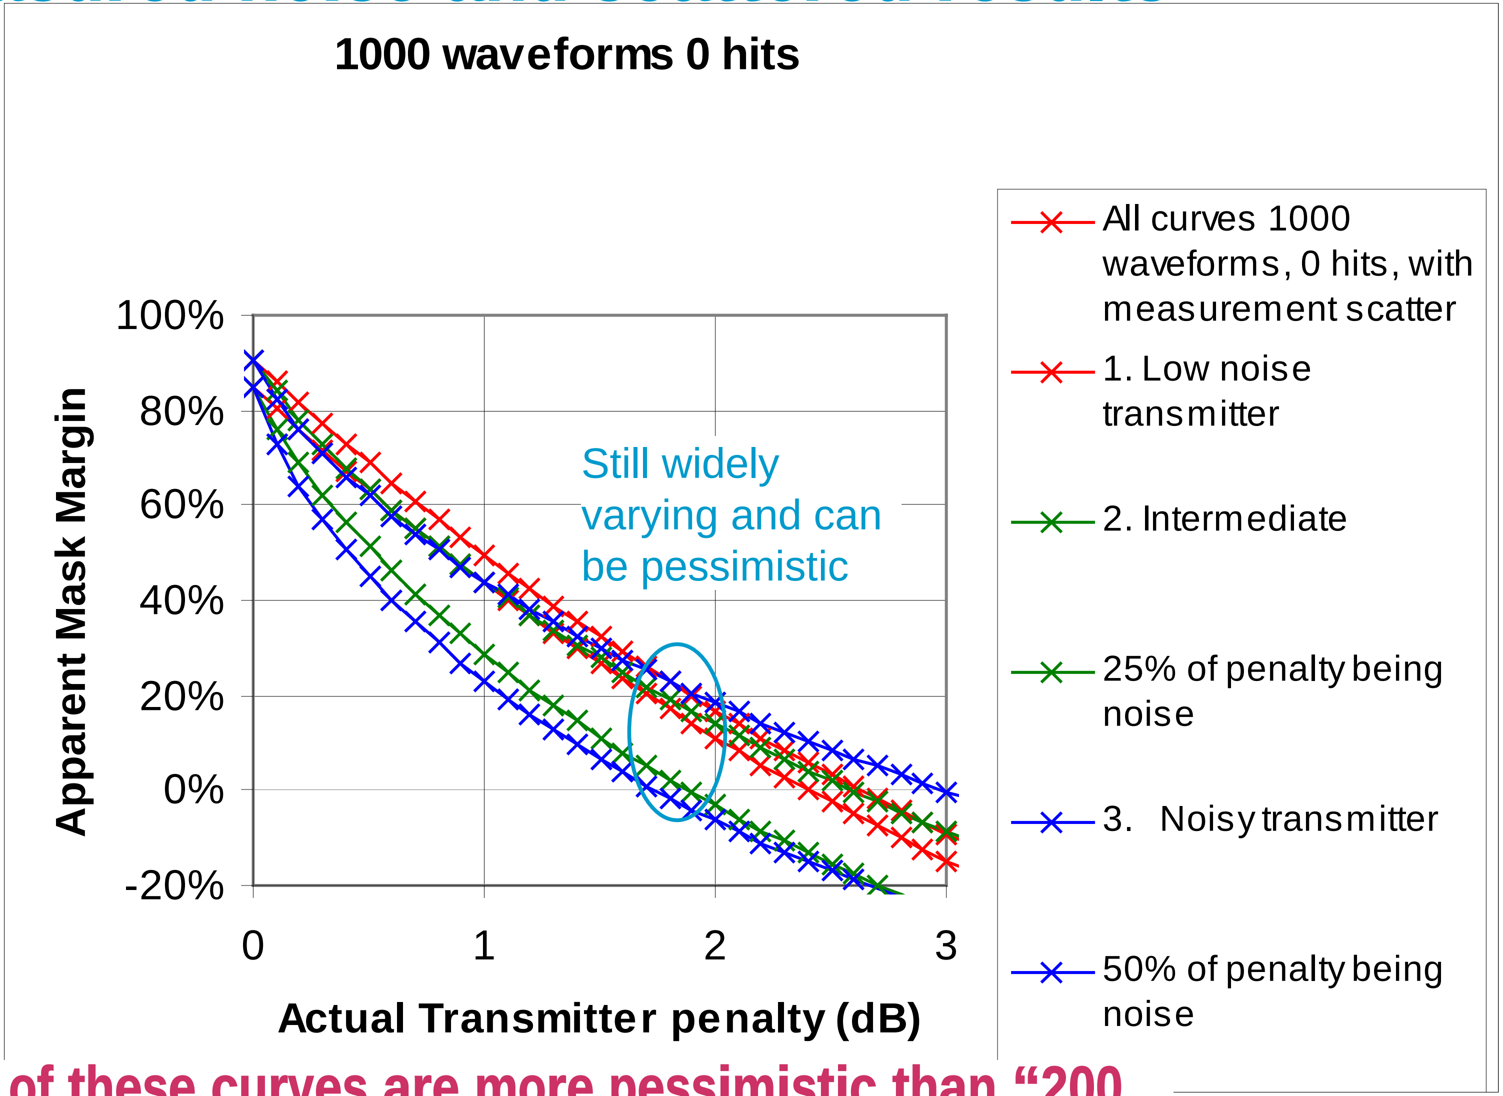

Some of these curves are more pessimistic than "200 waveforms" by up to 5%

Proposal for best accuracy

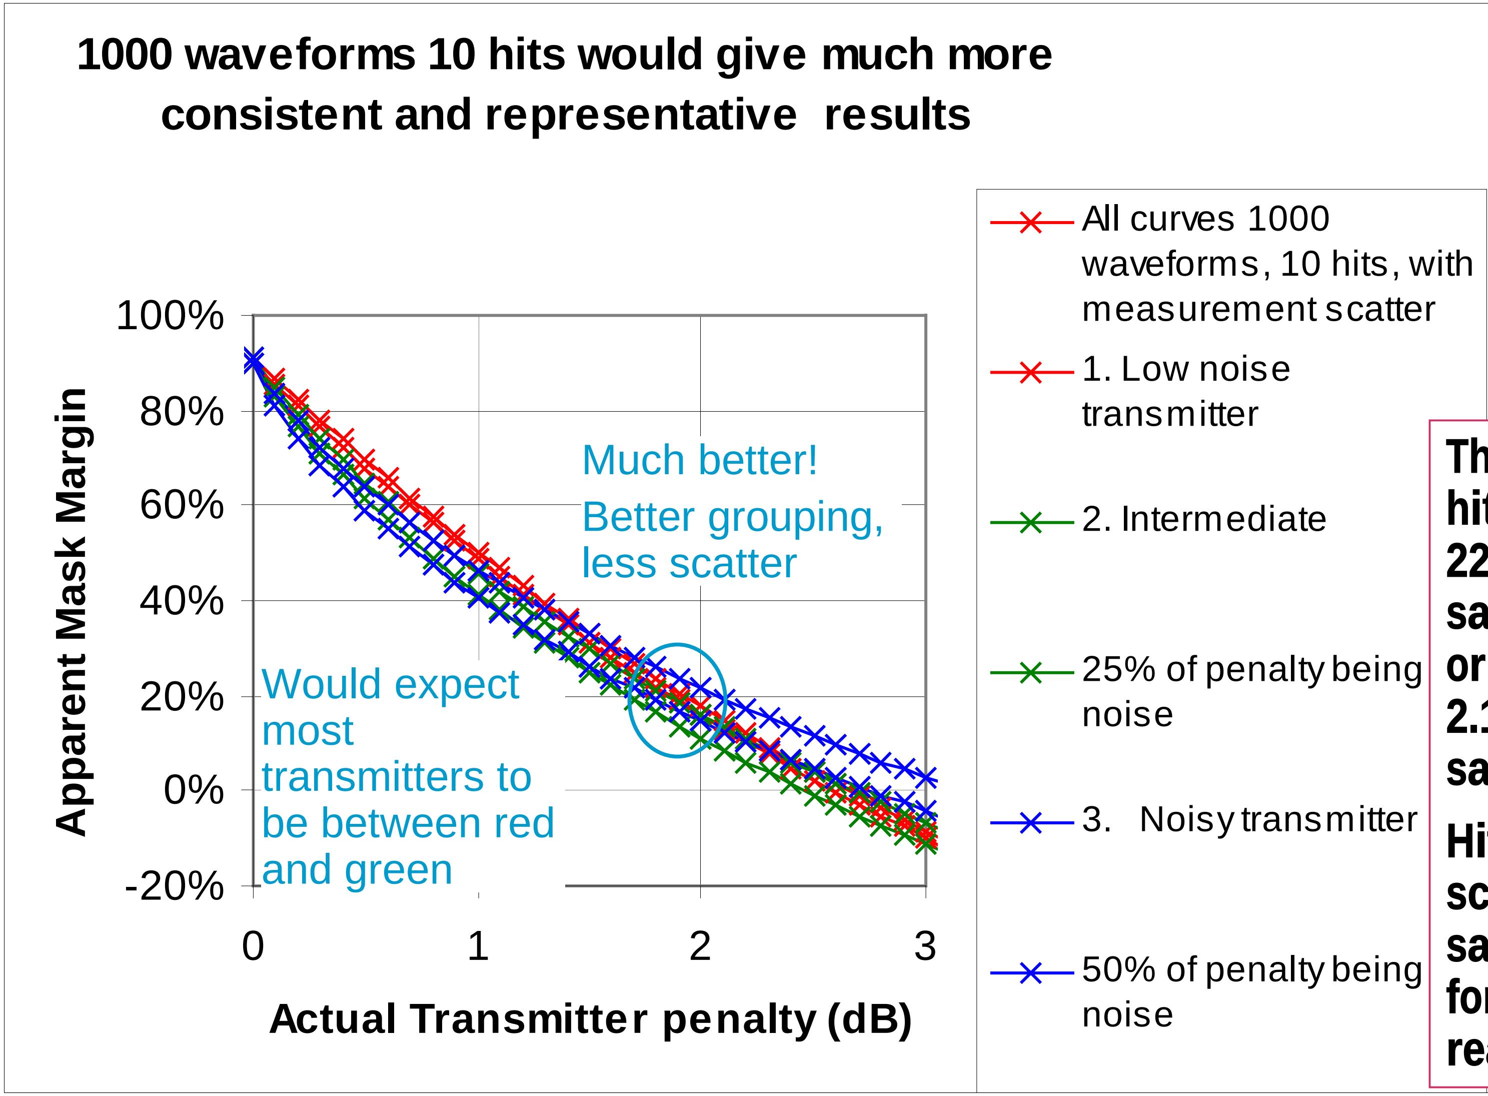

This is 10 hits in 225500 samples / UI or~1 hit per 2.10^4 samples / UI

Hit limit scales with sample size for stable reading

Numbers of waveforms, number of hits

Need a reasonable number of DCA “waveforms” to get samples from all likely patterns

e.g. 10 sweeps would be too few

We believe “200 waveforms” is industry baseline, and different companies add margin to this in different ways

This number is widely assumed but we can’t find it written down in any standard

Criterion of 0 hits gives poor reproducibility, for any number of waveforms

Increasing the number of samples with 0 hits over-measures random effects and may be affected by scope noise

And see backup

Summary

Mask margin based on very few or zero hits can be misleading and generate unrepeatable results

Increasing the number of waveforms reduces the apparent mask margin, does not make the measurement much more repeatable and over-measures transmitter and scope noise

Increasing the number of waveforms AND allowing a finite number of hits improves both measurement repeatability and relevance

Need to define a set proportion of hits / sample / UI

Too low a proportion: large measurement errors and unrepresentative (too pessimistic)

Too high a proportion: too lenient in some cases

1 in 2.10^4 hits / sample / UI: about right

10 sided masks in Ethernet standards should be adequate; no need to add more margin except for ageing

TDP test is the “gold standard”, not the mask

TDP measurement is designed to be representative of actual performance

Backup

SONET

Mask margin measurements with 4 sided masks (OC-48, OC-192) are expected to be less reproducible unless care is taken, because the convolution of a slanting trajectory of the waveform against the rectangular mask gives a “fat tailed distribution”

Relation of waveforms to samples

I have assumed 451 samples across scope screen and 2 UI / screen (20 ps / div). I believe that older scopes contemporary with OFSTP-4 would have been in this ballpark

“Number of waveforms” means number of samples x samples / screen

Thus “200 waveforms” gives 90200 samples in all, of which 45100 are in the eye compared with the mask

Modern scopes have controllable samples / screen up to 4096

Of the samples in the right UI (where the mask is), only 1 / 4 relate to 010 and 101 patterns which determine the penalty

Of these, only a fraction come anywhere near the mask (mask is 0.5 UI long anyway)

Signal and noise levels

signal to noise ratio = Q = (1 / 2 eye height) / (RMS noise)